SEMI-MARKOV PROCESSES: A PRIMER

Bennett Fox

August 1968

P-3577-2

AD 673 152

SEMI-MARKOV PROCESSES: A PRIMER

Bennett Fox

RAND Corporation
Santa Monica, California

August 1968

SEMI-MARKOV PROCESSES: A PRIMER

Bennett Fox*

The RAND Corporation, Santa Monica, California

Marrying renewal processes and Markov chains yields semi-Markov processes, and the former are special cases of the latter. In this expository paper, some of the main properties of the union are outlined.

*Any views expressed in this paper are those of the author. They should not be interpreted as reflecting the views of The RAND Corporation or the official opinion or policy of any of its governmental or private research sponsors. Papers are reproduced by The RAND Corporation as a courtesy to members of its staff.

Invited paper prepared for submission to Management Science.

I. INTRODUCTION

A key concept, already familiar to readers acquainted with queueing theory, is that of an imbedded Markov chain connecting regeneration points of a stochastic process. Starting from a regeneration point, the future is stochastically independent of the past. In the imbedded Markov chain, the original time scale of the transitions between regeneration points is replaced by a discrete time version where all transitions take unit time. The corresponding imbedded semi-Markov process looks only at the "states" corresponding to these regeneration points, but it does so in continuous time.

The use of these distinguished states, generally referred to as events by probabilists, is natural in a dynamic programming framework; see, e.g., Denardo and Mitten [13]. In fact, it was programming over semi-Markov processes that motivated our interest in these processes. We shall have more to say about this in Sec. 10.

Part of the definition of a distinguished state is its association with certain regeneration points. To fix this idea concretely, consider the M/G/1 queue (Poisson arrivals, general service time distribution, single channel). For many purposes, a convenient set of distinguished states is $\{0, 1, 2, \ldots\}$ where state i signifies i customers in the system and a service has just been completed.

This procedure is definitely in conflict with the notion of a state used in fully rigorous treatments of probability theory and semi-Markov processes in particular. There, a sample path of a stochastic process is defined as a function $X(\cdot, \omega): [0, \infty) \to E$ and the members of E are called states. This often necessitates speaking of a "holding

time" in a state. If this "holding time" is not necessarily exponentially distributed, as in semi-Markov processes, then we intuitively feel that the process does not remain in the same state during the holding period because the process has memory (is not Markovian). By contrast, all our distinguished states are occupied only for an instant. This procedure is unconventional and heuristic, but this is not supposed to be a review paper for experts. Rather, it is directed toward those readers with some prior exposure to Markov chains and renewal theory who would like to get a feel for what semi-Markov processes are all about and how they arise in applications. For many readers belonging to this class, it is felt that this primer is a more accessible first introduction than the original papers cited in the reference list.

Letting $N_j(t)$ denote the number of times state j is entered in the half-open interval (0, t], we obtain the Markov Renewal Process (MRP) $\underline{N}(t) = (N_o(t), N_1(t), N_2(t), \ldots)$. In the M/G/1 queue, for example, $N_o(t)$ is the number of busy periods completed in (0, t].

Let Z(t), the semi-Markov process, be the last distinguished state entered in [0, t]. In general, such a last state is not well defined, but in the applications there is virtually never any difficulty. See Sec. 3 for discussion of this point. In the M/G/1 queue, if at the last service completion epoch in [0, t] there were i customers in the system, then Z(t) = i; however, at time t there may be more than i customers in the system due to arrivals since the last service completion.

I^+ is defined to be the set of distinguished states, assumed countable. The state transitions form a Markov chain with transition probabilities (p_{ij}), where direct transitions from a state to itself (e.g., $p_{ii} > 0$) are allowed. Given that an $i \to j$ transition is about

to occur, the duration of the transition has distribution F_{ij}. In the M/G/1 queue, $F_{ij} = G$ if $i > 0$ and F_{oj} is the convolution of G with the exponential distribution whose mean is the reciprocal of the arrival rate; the transition matrix is given in Sec. 7.

Semi-Markov processes (SMP's), first studied by P. Lévy and W. L. Smith, generalize several familiar processes. Note that

(i) a one-state SMP is a renewal process;

(ii) an SMP with F_{ij} degenerate at one for all i, j is a Markov chain;

(iii) an SMP with all F_{ij} exponential and independent of j is a continuous time countable-state Markov process;

(iv) an alternating renewal process is a two-state SMP.

All SMP's have renewal processes imbedded within them corresponding to looking only at successive returns to the same state. In the M/G/1 queue, if the state in question is 0, then we are looking at successive returns to the beginning of an idle period.

Many problems in management science and operations research can be modeled as SMP's: for example, queueing, inventory, and maintenance problems. Explicit recognition of the underlying SMP often streamlines the analysis. For details, see, e.g., Pyke [44], Barlow and Proschan [1], Fabens [15], Foley [18], Çinlar [4-8], and Neuts [37-41]. These areas by no means exhaust the possibilities; e.g., Perrin and Sheps [43] and Weiss and Zelen [52] apply SMP's to medical problems. John McCall (unpublished work) has used SMP's to model movements among income classes in a study of strategies for combatting poverty. For proofs, citations of earlier papers, and additional topics in SMP's, the reader should consult the reference list. Another expository paper is Janssen [26].

Recently, Neuts [42] has published a bibliography on SMP's. I thank Professor Neuts for bringing several relevant papers to my attention.

II. DIST.:IGUISHED STATES

In this section we examine precisely what is meant by a distinguished state. Since discussing this topic in an offhand manner could result in confusion, the subject is treated in some detail.

For each sample path of a stochastic process X, there is a correspondence between $[t: \ t \geq 0]$ and a set of states S. If every state (in our sense of the word) in S is required to have the Markov property, then in general S will be uncountable since a history of the process, or at least the relevant portion of it, must be part of the state definition. However, all that we require of S is that it have an appropriate countable subset I^+ of (distinguished) states having the Markov property. Thus, for a state to be a candidate for I^+, it must correspond to a regeneration point, but we do not require that all states corresponding to regeneration points belong to I^+. We assume that the process starts in a distinguished state at time 0.

Note that the set of distinguished states used previously in our discussion of the M/G/1 queue does not include all regeneration points, since any time the system is idle (empty), it is at a regeneration point. Occasionally it is convenient to add to I^+ the state corresponding to arrival epochs to an empty system. Our choice of distinguished states conforms to our requirements because the time to the next arrival is stochastically independent of the time elapsed since the last arrival. In general, arrival epochs, except those corresponding to the start of a busy period, are not regeneration points. Thus the state "i customers in the system and a customer has just arrived" cannot be a distinguished state, unless

$i = 0$ or the service times are exponential. Through the use of so-called supplementary variables, we can define every state in the original process so that it is Markovian. Each state is then a couple of the form (i,u), which denotes i customers in the system and the customer being processed has been in service for time u; for $i = 0$, u is arbitrary, say 0. Sometimes supplementary variable techniques are useful as an alternative or adjunct to SMP techniques; see, e.g., Cox and Miller [9]. A disadvantage of supplementary variable techniques is that superfluous regularity conditions, e.g., absolutely continuous transition time distributions, often must be imposed on the original process to justify their use.

Returning to the general discussion, we require the distinguished states to be defined such that non-zero holding times in a distinguished state are forbidden[1] but instantaneous transitions among the distinguished states are allowed. This is a departure from the setup of Pyke [44], although the two formulations are essentially equivalent. The notion of an auxiliary path (see [47]), needed in the conventional setup to handle such processes as the M/G/1 queue, is not required by us. Our definition of distinguished state permits a graphic representation of SMP's in terms of networks with branch nodes (distinguished states) and stochastic arc lengths; see [19]. For example, traversing an arc could correspond to a customer completing service.

---

[1] By contrast, in the conventional setup non-zero holding times are essential, but this seems to be an artifice unless at every instant the process is memoryless, i.e., Markovian.

To remove ambiguity in case of instantaneous transitions,[2] we define $X^{+}(t) = X(t^{+})$; thus, X^{+} is right continuous and the last distinguished state of X^{+} entered in $[0, t]$, say, is well defined, provided that the process does not explode; see the discussion of regularity in Sec. 3. Since we have prohibited non-zero holding times in distinguished states, we cannot allow a distinguished state to correspond to a nondegenerate interval of regeneration points (e.g., an idle period in an M/G/1 queue). Thus, to exclude an infinite sequence of instantaneous transitions from a state to itself, we require that the distinguished states be defined such that, for all nondegenerate intervals (a, b), $i \in I^{+} \Rightarrow P\{X(t) = i, \forall t \in (a, b)\} = 0$. For example, in the M/G/1 queue it does not suffice to define the distinguished state 0 as 0 customers in the system. The condition that a service has just been completed must be added.

In applications, the first step is to carefully specify the distinguished states, which throughout the sequel are simply called "states," necessitating definitions slightly different from conventional usage.

---

[2] Sometimes it is convenient to permit instantaneous transitions; see, e.g., Denardo [11]. Yackel [55] makes a detailed study of limit theorems for SMP's with instantaneous transitions allowed.

III. REGULAR SMP'S

In the literature the Z process is called an SMP. However, the MRP and the SMP are different aspects of the same underlying stochastic process; therefore, by slight abuse of language, we shall refer to the underlying process itself as an SMP or a MRP, using the terms interchangeably.

An MRP is <u>regular</u> if with probability one (w.p.1) each state is entered only a finite number of times in any finite time span--i.e., if $P[N_i(t) < \infty] = 1$, $\forall i \in I^+$ and $t \geq 0$. An MRP is <u>strongly regular</u> if w.p.1 the total number of state transitions is finite in any finite time span--i.e., if $P[\Sigma N_i(t) < \infty] = 1$, $\forall t \geq 0$. Clearly strong regularity implies regularity and, if $n < \infty$, it suffices that $\underline{H} = (H_o, \ldots, H_n)$ have at least one component nondegenerate at zero for every ergodic subchain of the imbedded Markov chain, where H_i is the unconditional distribution of time elapsed starting from state i until the next state is entered (possibly i itself). In the denumerable state case $(n = \infty)$, see Pyke [44] and Pyke and Schaufele [46] for conditions that imply strong regularity; see also Feller [17]. In the sequel, we assume that strong regularity holds.

IV. FIRST PASSAGE AND COUNTING DISTRIBUTIONS

Let[3]

$$Q_{ij}(t) = p_{ij}F_{ij}(t)$$

$$H_i(t) = \sum_j Q_{ij}(t)$$

$$P_{ij}(t) = P[Z(t) = j|Z(0) = i]$$

$$G_{ij}(t) = P[N_j(t) > 0|Z(0) = i]$$

(first passage time distribution)

$$M_{ij}(t) = E[N_j(t)|Z(0) = i]$$

(mean entry counting function).

Defining the convolution

$$(A * B)(t) = \int_{0^-}^{t} A(t - x)\, dB(x)$$

and deleting the argument t below, we have by straightforward renewal-theoretic arguments:

$$P_{ij} = (1 - H_i)\delta_{ij} + \sum_k Q_{ik} * P_{kj} = (1 - H_i)\delta_{ij} + P_{jj} * G_{ij}$$

$$G_{ij} = Q_{ij} + \sum_{k \neq j} Q_{ik} * G_{kj}$$

$$M_{ij} = G_{ij} + G_{ij} * M_{jj} = Q_{ij} + \sum_k Q_{ik} * M_{kj}.$$

[3]Unless otherwise stated, all summations will be over I^+ and all functions vanish for negative arguments.

In general, these relations cannot be solved analytically, although for the moments complicated expressions have been obtained (Pyke and Schaufele [46]). In the finite state case, numerical solutions can be obtained by numerical inversion of the corresponding Laplace transforms (see, e.g., [2], [24], [33], and [47]). For each value of s, only one matrix inversion in the transform domain is required--that of $I - q(s)$, where $s > 0$ and

$$q(s) = \left(\int_{0^-}^{\infty} e^{-st}\, dQ_{ij}(t)\right)$$

In obvious notation, having found $[I - q(s)]^{-1}$, either analytically as a function of s or, for suitably spaced values of s, numerically, one successively computes

$$m(s) = [I - q(s)]^{-1}q(s) = [I - q(s)]^{-1} - I$$

$$g_{ij}(s) = m_{ij}(s)/[1 + m_{jj}(s)]$$

$$p_{ij}(s) = p_{jj}(s)g_{ij}(s), \quad i \neq j$$

$$p_{jj}(s) = \frac{1 - h_j(s)}{1 - g_{jj}(s)} ,$$

and then inverts the transforms. Although this procedure is not trivial, it often compares favorably with the alternative simulation approach for getting the transient behavior in the time domain. By usual limit theorems for Laplace transforms (Widder [53], Feller [16], see also Jewell [27]), the behavior in the time domain for large (small) t corresponds to behavior in the transform domain for small (large) s. Renewal theory provides an important tool in studying asymptotic behavior; see Sec. 8. For the stationary probabilities, see Sec. 6.

Conditioning on the event that no state in a subset B of I^+ is entered in (0, t] may be of interest. For example,

$${}_B P_{ij}(t) = P[Z(t) = j \mid Z(0) = i,\ N_k(t) = 0,\ \forall k \in B]$$

$${}_B G_{ij}(t) = P[N_j(t) > 0 \mid Z(0) = i,\ N_k(t) = 0,\ \forall k \in B]$$

$${}_B M_{ij}(t) = E[N_j(t) \mid Z(0) = i,\ N_k(t) = 0,\ \forall k \in B]$$

can be calculated from the formulas already given by (temporarily) making the states in B absorbing.

Barlow and Proschan [1, pp. 132-134], using "renewal" arguments, show that the first and second moments of G_{ij}, denoted respectively by μ_{ij} and $\mu_{ij}^{(2)}$, are given by

$$\mu_{ij} = \sum_{k \neq j} p_{ik}\mu_{kj} + \nu_i$$

$$\mu_{ij}^{(2)} = \sum_{k \neq j} p_{ik}\left[\mu_{kj}^{(2)} + 2\nu_{ik}\mu_{kj}\right] + \nu_i^{(2)} ,$$

where ν_{ij} is the mean of F_{ij}, $\nu^{(1)} = \nu$, and

$$\nu_i^{(\ell)} = \int_{0^-}^{\infty} t^{\ell}\, dH_i(t).$$

We assume that $\nu_i^{(2)} < \infty$, $\forall i \in I^+$. If the imbedded Markov chain is finite and ergodic, these equations have a unique finite solution[4] and, with π_j the stationary probability that the last state entered

[4]See appendices 1 and 2 of Fox [19] for an efficient way to solve these equations. (An expression for the "bias terms" in Markov renewal programming involves the first passage time moments, which are of intrinsic interest, but recently Jewell [29] has derived a remarkably simple alternative expression, obviating the need to calculate $\{\mu_{ij}^{(2)}\}$ to evaluate the bias terms.)

is j if all F_{ij} were degenerate at one,[5] multiplying μ_{ij} and $\mu_{ij}^{(2)}$ by π_i and summing yields

$$\mu_{jj} = (1/\pi_j) \sum_k \pi_k \nu_k$$

$$\mu_{jj}^{(2)} = (1/\pi_j)\left[\sum \pi_k \nu_k^{(2)} + 2 \sum_{k \neq j} \sum_i \pi_i P_{ik} \nu_{ik} \mu_{kj}\right].$$

For finite state SMP's, the probability that state j is ultimately reached starting from i is[6]

$$G_{ij}(\infty) = \begin{cases} 1, & \text{if } i, j \in E_k \\ 0, & \text{if } i \in E_k,\ j \in E_\ell,\ k \neq j \\ [(I - A)^{-1}\theta]_i, & \text{if } i \in T,\ j \in E_k \end{cases}$$

where A is the submatrix of P corresponding to the set T of transient states, $E_1, \ldots, E_m$ are the recurrent subchains of P, and

$$\theta_i = \sum_{\ell \in E_k} P_{i\ell}, \qquad i \in T.$$

The mean time to leave T starting from i is

$$\xi_i = [(I - A)^{-1} \nu^t]_i, \qquad i \in T,$$

where ν^t is the vector of ν_j's, $j \in T$. Pyke [45] obtains a double generating function for the distribution of $N_j(t)$, viz.,

[5] In other words, π is the stationary measure for the imbedded chain, but not (in general) for the SMP itself.

[6] The case $i, j \in T$, of less interest, is not considered.

$$\Psi_z = \underline{1} - (1 - z)m\,[z\,I + (1 - z)D]^{-1},$$

where $\underline{1}$ is a matrix of 1's,

$$D = \left((1 - q_{ij})\delta_{ij}\right)^{-1} = \left(\frac{\delta_{ij}}{1 - q_{ij}}\right)$$

and

$$\Psi_z = (\phi_{ij}(z;\,\bullet))$$

$$\phi_{ij}(z;\,s) = \int_{0^-}^{\infty} e^{-st}\, d_t w_{ij}(z;\,t)$$

$$w_{ij}(z;\,t) = \sum_{k=0}^{\infty} z^k v_{ij}(k;\,t)$$

$$v_{ij}(k;\,t) = P[N_j(t) = k \,|\, Z(0) = i].$$

Thus, in principle, the probabilities and moments can be obtained in the usual way. The Laplace transform m of the first moment $(M_{ij}(\bullet))$ was already given. See Pyke and Schaufele [46] for further general moment computations, weak and strong laws of large numbers, and central limit theorems. A generating function that yields many quantities of interest upon considering special cases has been obtained by Neuts [36]. Stone [50] derives the distribution of the maximum of an SMP.

V. STATE CLASSIFICATION

In classifying the states of an SMP transient, null recurrent, or positive recurrent, we must distinguish between a state's classification in the imbedded Markov chain and in the SMP itself. For I^+ finite and $\nu_i < \infty$, $\forall i \in I^+$, the distinction disappears and a state j is either transient or positive recurrent ($G_{jj}(\infty) = 1$ and $\mu_{jj} < \infty$). In large-scale applications, the ergodic subchain-transient set breakdown may not be obvious and recourse may be necessary to an algorithmic classification scheme such as that of Fox and Landi [23].

For I^+ infinite, a state j is transient (recurrent--i.e., $G_{jj}(\infty) = 1$) in the SMP $\Leftrightarrow$ j is transient (recurrent) in the imbedded Markov chain. State ℓ is positive recurrent in the imbedded Markov chain (contained in ergodic subchain E_k) and, for some constant c, $\nu_{ij} \le c < \infty$, $\forall i$, $j \in E_k$, $\Rightarrow$ ℓ is positive recurrent in the SMP. An SMP is positive recurrent if all the states in I^+ are positive recurrent in the SMP.

We remark that, if $Z(0) = i$ and i belongs to the same positive recurrent ergodic subchain as j, $t^{-1}N_j(t) \to 1/\mu_{jj}$ w.p.1, a strong law that follows immediately from renewal theory. Under these conditions and assuming $\mu_{jj}^{(2)} < \infty$, $N_j(t)$ is asymptotically normally distributed with mean t/μ_{jj} and variance $t\mu_{jj}^{-3}(\mu_{jj}^{(2)} - \mu_{jj}^2)$, a consequence of a renewal-theoretic result found, for example, in Feller [16, p. 359].

VI. STATIONARY PROBABILITIES

It is important to distinguish between the stationary probabilities $\{\pi_i\}$ with respect to the imbedded Markov chain and the stationary probabilities $\{\rho_i\}$ with respect to the SMP.[7] Thus ρ_i is the steady state probability that the last distinguished state entered is i. Hence the $\{\rho_i\}$ are of direct interest in applications, while the $\{\pi_i\}$ are computed only as an intermediate step. We consider first the case I^+ finite and $\nu_i < \infty$, $\forall i \in I^+$.

$$\rho_j = \begin{cases} \nu_j/\mu_{jj}, & j \in E_k,\ Z(0) \in E_k \\ G_{ij}(\infty)\nu_j/\mu_{jj}, & j \in E_k,\ Z(0) = i \in T \\ 0, & j \in T \\ 0, & j \in E_k,\ Z(0) \in E_\ell,\ k \neq \ell \end{cases}$$

where $G_{ij}(\infty)$ was computed already and

$$\nu_j/\mu_{jj} = \frac{\pi_j \nu_j}{\sum_{i \in E_k} \pi_i \nu_i}$$

with $\{\pi_i\}$ here being the stationary probabilities for the imbedded Markov chain given that $Z(0) \in E_k$.

In the remainder of this section we assume that the imbedded Markov chain is irreducible and that the SMP is positive recurrent, where I^+ may be finite or infinite. We also assume that the mean

---

[7] In general, the stationary probabilities must be interpreted as Cesaro limits. If the process is aperiodic, these reduce to ordinary limits.

transition times are uniformly bounded away from zero; i.e., $0 < \epsilon \leq \nu_{ij} < \infty$. With these assumptions, Fabens [15] shows that

$$\rho_j = \frac{\pi_j \nu_j}{\sum_i \pi_i \nu_i}$$

in agreement with results given above for the I^+ finite case. Define $\sigma(x)$ as the time of the last transition completion before or at x and $\tau(x)$ as the time of the next transiton completion after x. The random variables

$$\gamma(x) = \tau(x) - x \quad \text{(excess r.v.)}$$

$$\delta(x) = x - \sigma(x) \quad \text{(shortage r.v.)}$$

are of interest. Adding to the previous assumptions the hypotheses that the mean recurrence times $\{\mu_{ii}\}$ are finite and that $Z(\bullet)$ is aperiodic, Fabens shows that

$$\lim_{t\to\infty} P\{\delta(t) \leq x | Z(t) = i\} = \lim_{t\to\infty} P\{\gamma(t) \leq x | Z(t) = i\}$$

$$= \frac{1}{\nu_i} \int_0^x [1 - H_i(u)]\, du.$$

This generalizes the well-known result from renewal theory for the one state case, obtained there as a corollary to the key renewal theorem; see, e.g., Barlow and Proschan [1].

The general question of existence and uniqueness of stationary measures is dealt with in Pyke and Schaufele [47]. Cheong [3] gives conditions under which convergence to the steady state is geometric;

see also Teugels [51]. An estimate of the convergence rate is important. If it is high enough, troublesome transient phenomena can be neglected. We then pass directly to a relatively simple steady state analysis.

VII. EXAMPLE: THE M/G/1 QUEUE

To illustrate the notion of stationary probabilities for an SMP, we consider the M/G/1 queue. Let[8]

π_i = the stationary probability that i customers are in the system just after a random service completion epoch

ρ_i = the stationary probability that i customers are in the system just after the service completion epoch preceding a random point in time

p_i = the stationary probability that i customers are in the system at a random point in time.

We assume that the traffic intensity λb is less than one, where λ is the arrival rate and b the mean service time, assumed positive. Although it is easily shown that

$$\rho_i = \begin{cases} \lambda b \pi_i, & i \geq 1 \\ (1 + \lambda b)\pi_0 = 1 - (\lambda b)^2, & i = 0, \end{cases}$$

it turns out that $p_i = \pi_i$, $\forall i$, a remarkable result originally due to Khintchine [31] and derived in a more elementary manner by Fox and Miller [24] using SMP theory. In bulk queues (Fabens [15]), for example, the stationary measures for the imbedded Markov chain and the original queueing process are different.

Readers familiar with queueing theory may prefer to skip to the last paragraph of this section. In between, the standard manipulations yielding $G_\pi(z)$, the generating function of the $\{\pi_i\}$, are performed.

---

[8] The stationary probabilities for the imbedded chain and the SMP are $\{\pi_i\}$ and $\{\rho_i\}$, respectively.

Recalling that state n means that there are n people in the system and a service has just been completed, we obtain the well known transition matrix for the imbedded Markov chain:

$$P = \begin{array}{c|cccccccc} & 0 & 1 & 2 & 3 & 4 & \cdot & \cdot & \cdot \\ \hline 0 & k_0 & k_1 & k_2 & k_3 & k_4 & \cdot & \cdot & \cdot \\ 1 & k_0 & k_1 & k_2 & k_3 & k_4 & \cdot & \cdot & \cdot \\ 2 & 0 & k_0 & k_1 & k_2 & k_3 & \cdot & \cdot & \cdot \\ 3 & 0 & 0 & k_0 & k_1 & k_2 & \cdot & \cdot & \cdot \\ 4 & 0 & 0 & 0 & k_0 & k_1 & & & \\ \cdot & \cdot & \cdot & \cdot & \cdot & & \cdot & & \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & & \cdot & \\ \cdot & \cdot & \cdot & \cdot & \cdot & & & & \cdot \end{array}$$

where the probability that n customers arrive while a customer is being serviced is

$$k_n = \int_{0^-}^{\infty} \frac{e^{-\lambda t}}{n!} (\lambda t)^n \, dB(t)$$

and B is the service distribution. By the usual straightforward manipulations, we find that the generating function of the $\{k_n\}$ is

$$G_k(z) = \sum_i k_i z^i = \beta(\lambda(1 - z)),$$

where β is the Laplace-Stieltjes transform of B, i.e.,

$$\beta(s) = \int_{0^-}^{\infty} e^{-st} \, dB(t).$$

To obtain the stationary vector π for the chain, we multiply the i-th relation determined by $\pi P = \pi$ by z^i and sum, define the generating function

$$G_\pi(z) = \sum_i \pi_i z^i,$$

and obtain from the special form of P for this chain by an easy calculation the standard result

$$G_\pi(z) = \frac{\pi_o(1 - z)G_k(z)}{G_k(z) - z} .$$

Using the fact that $\lim_{z\to 1^-} G_\pi(z) = 1$ (i.e., the probabilities sum to 1) and applying L'Hospital's rule,

$$\pi_o = 1 - \lambda b.$$

Summarizing our results so far,

$$G_\pi(z) = \frac{\pi_o(1 - z)\beta(\lambda(1 - z))}{\beta(\lambda(1 - z)) - z} .$$

Thus the mean number in the system averaged over service completion epochs is, with σ^2 the variance of the service times,

$$\lim_{z\to 1^-} G_\pi(z) = \pi_o \lim_{z\to 1^-} \left[\frac{G''_k(z) - 2G'_k(z)(G'_k(z) - 1)}{2(G'_k(z) - 1)^2}\right] = \frac{(\lambda b)^2 + \sigma^2}{1 - \lambda b} + \lambda b,$$

and by the fact that $\pi_i = p_i$, $\forall i$, is also the mean number in the system at random point in time (in the steady state). Higher moments

and probabilities can be obtained from the generating function by appropriate differentiations, which, however, become quite tedious.

Having found $G_\pi(z)$, the Laplace-Stieltjes transform of the stationary waiting distribution (Pollaczek-Khintchine formula) for the first come, first served (FIFO) discipline can easily be found. The derivation depends on the fact that, since the arrival process is Poisson, an arrival plays the role of a random observer. If an arrival finds the system empty, the conditional wait in queue is 0. Otherwise it is governed by the remaining processing time of the customer in service, the excess random variable, plus the service times for the customers (if any) already in queue. Noting that $G_p(z) = G_\pi(z)$, the interested reader can readily derive a version of the Pollaczek-Khintchine formula, namely, $\pi_o + \frac{1 - \beta(s)}{sb\beta(s)} [G_p(\beta(s)) - \pi_o]$. See Feller [16, p. 392], for an alternate elegant derivation that bypasses the calculation of $G_p(z)$. A third derivation follows from the fact that the number of customers in the system just after a departure is the number of arrivals during his total wait (queueing time plus service time); the resulting equation is solved by taking generating functions yielding the standard form of the Laplace-Stieltjes transform of the stationary queueing delay distribution $s\pi_o/(s - \lambda(1 - \beta(s)))$, the more familiar version of the Pollaczek-Khintchine formula.[9] For a fourth derivation, where the (superfluous) assumption of an absolutely continuous failure distribution is tacitly made, see Cox and Miller [9, pp. 241-242].

[9] Comparison of the two versions yields an interesting and surprising identity.

VIII. ASYMPTOTIC FORM OF M_{ij}

From a previous section, we know that

$$m_{ij}(s) = g_{ij}(s)[1 + m_{jj}(s)]$$

$$= g_{ij}(s)\left[1 + \frac{g_{jj}(s)}{1 - g_{jj}(s)}\right]$$

Formally expanding e^{-sx} in a Taylor series under the integral defining g_{ij}, integrating termwise, and performing the appropriate algebraic manipulations yields for i, j in the same ergodic subchain

$$m_{ij}(s) = \frac{1}{s\mu_{jj}} + \frac{\mu_{jj}^{(2)}}{2\mu_{jj}^{2}} - \frac{\mu_{ij}}{\mu_{jj}} + o(1),$$

whence by a Tauberian argument

$$M_{ij}(t) - \frac{t}{\mu_{jj}} \xrightarrow[\text{(Cesàro)}]{} \frac{\mu_{jj}^{(2)}}{2\mu_{jj}^{2}} - \frac{\mu_{ij}}{\mu_{jj}},$$

a result that can be obtained by analogy with renewal theory for delayed recurrent events, where the time to the first "renewal" has distribution G_{ij} and the spacing between subsequent renewals has distribution G_{jj}. If the SMP is aperiodic, the Cesàro limit reduces to an ordinary limit. It can be shown that the formal manipulation used to obtain the asymptotic expansion of $m_{ij}(s)$ is justified if $\mu_{jj}^{(2)} < \infty$. If I^{+} is finite, $\nu_{i}^{(2)} < \infty$, $\forall i \in I^{+} \Rightarrow \mu_{jj}^{(2)} < \infty$. A result

that holds for all t is $Z(0) = j \Rightarrow M_{jj}(t) \geq t/\mu_{jj} - 1$, which follows from Barlow and Proschan [1, Theorem 2.5]. We can obtain a tighter inequality from [1] and, if G_{jj} has increasing failure rate, an upper bound as well.

IX. FINITE SMP'S WITH COSTS

Sometimes the performance of a system is evaluated by probabilistic criteria, such as average delay in a queue, that serve as surrogates for monetary loss. It is more appealing to deal directly with expected loss as a performance measure.[10] In this section we indicate how to do this.

Often in applications, costs are associated with the transitions. Measuring time from the start of an $i \to j$ transition, let $C_{ij}(x|t)$ be the cost incurred up to time x given that the transition length is t. The expected discounted cost for a transition starting from state i is then

$$\gamma_i(\alpha) = \sum_j p_{ij} \int_{0^-}^{\infty} dF_{ij}(t) \int_{0^-}^{t} e^{-\alpha x} \, d_x C_{ij}(x|t),$$

where a cost incurred at time x is discounted by the factor $e^{-\alpha x}$ An elementary renewal type argument then shows that $v_i(\alpha)$, the total expected discounted cost over an infinite horizon starting from state i, satisfies

$$v(\alpha) = \gamma(\alpha) + q(\alpha)v(\alpha),$$

where $\alpha > 0$ and $v(\alpha)$ and $\gamma(\alpha)$ are the vectors with components $v_i(\alpha)$ and $\gamma(\alpha)$, respectively. Thus, assuming a finite number of states,

$$v(\alpha) = [I - q(\alpha)]^{-1}\gamma(\alpha).$$

[10] A similar remark applies to "chance-constrained" programming.

Since I - q(0) is singular and a direct asymptotic expansion is not obvious, it is convenient to make use of the relation between q and m given earlier to study the behavior of $v(\alpha)$ as $\alpha \to 0^+$. Following Jewell [27], we have

$$v(\alpha) = [I + m(\alpha)]\gamma(\alpha).$$

Making use of the expansion of $m(\alpha)$ given in the preceding section, we find that, if i is a recurrent state, $v_i(\alpha)$ has the form

$$v_i(\alpha) = \ell_i/\alpha + w_i + o(1)$$

and a straightforward argument in Fox [19] then shows that this form is valid for any state; i.e.,

$$v(\alpha) = \ell/\alpha + w + o(1) ,$$

where expressions for ℓ, the loss rate vector, and w, the bias term vector, can be found in Jewell [28, 29] and Fox [19], where appropriate conditions are given to justify the expansion. Substituting this relation into $v(\alpha) = \gamma(\alpha) + q(\alpha)v(\alpha)$ and equating the coefficients of α^{-1} and the constant terms, respectively, yields[11]

$$P\ell = \ell$$

$$\gamma + Pw = w + y$$

$$y_i = \sum_j p_{ij}\nu_{ij}\ell_j$$

$$= \nu_i\ell_i, \quad \text{if } i \text{ is recurrent.}$$

[11]This procedure can be justified by a simple contradiction argument. Note that $q_{ij}(\alpha) = p_{ij}(1 - \alpha\nu_{ij}) + o(\alpha)$ and that the loss rate for all states in an ergodic subchain is the same.

These expressions can be solved uniquely for ℓ, but w is determined only up to an additive constant in each ergodic subchain; see, e.g., Denardo and Fox [12]. An interesting and intuitive result that follows easily from the above formulas is that the loss rate for each state in an ergodic subchain E_k is the same and equal to $\pi^{(k)}\gamma'/\pi^{(k)}\nu'$, where $\pi^{(k)}$ is the stationary vector for the corresponding submatrix and γ' and ν' are the restrictions γ and ν, respectively, to E_k. This formula can be rewritten as

$$\sum_{i\in E_k} \rho_i(\gamma_i/\nu_i),$$

which is the sum of the expected cost per unit time in each state of E_k weighted by the respective stationary probabilities for the SMP. The loss rates for the transient states are obtained from the fact reflected in $P\ell = \ell$ that the loss rate for a state is given by the appropriate convex combination and that I - A, where A is the submatrix corresponding to the transient states, is invertible.

Denoting the undiscounted loss up to time t by L(t), we obtain from the asymptotic expansion of $v(\alpha)$ that

$$L(t) - \ell t \xrightarrow[\text{(Cesàro)}]{} w.$$

Jewell [30] studies the fluctuations in cumulative loss in what is essentially the one-state case. If the imbedded Markov chain is ergodic, these results extend in principle to n-state problems by considering G_{11} and the distribution of cumulative loss until the first return. In general, the calculation would be tedious. Besides,

we are distinguishing here only one state out of n. Apparently, no one has dealt with fluctuation theory for the n-state case directly.

A related topic is a central limit theorem for cumulative loss. Since cumulative loss is an example of a functional of a Markov Renewal Process, results of Pyke and Schaufele [46] apply.

X. MARKOV RENEWAL PROGRAMMING

The situation becomes more interesting when, at each state i, one has a set of options A_i and the choice at i simultaneously determines p_{ij}, F_{ij}, and C_{ij} for all $j \in I^+$. The goal is to choose a policy that minimizes either the expected discounted loss or the loss rate. In the latter case, an appropriate secondary objective is to minimize (ℓ, w) lexicographically, which is especially important when some policies can have transient states. With either criterion, an optimal policy can be found by linear programming when I^+ and $\times_{i \in I^+} A_i$ are finite. For details, see, e.g., Jewell [28, 29], Fox [19, 22], Denardo [10, 11], and Denardo and Fox [12], where references to the earlier (extensive) literature on the subject are given. The linear programming formulation facilitates sensitivity analyses and parametric studies. Controlling roundoff errors is probably less difficult in the averaging version; see related remarks in [12].

Some papers treat the I^+ infinite case, but the author believes that, for applications, the general theory developed so far for that case is inadequate and that particular problems are best attacked on an ad hoc basis. The averaging version of the infinite state case apparently has been studied only in the discrete time setup; see, e.g., Derman [14], Ross [48]. However, in the discounted continuous time version no new theoretical problems arise when the problem is approached via contraction mappings (Denardo [10]). For the case where I^+ is finite but the finiteness restriction on $\times A_i$ is dropped, see Fox [20]. The connection with generalized linear programming and column generators is outlined in [12].

Miller [32, 33] treats the continuous time Markov process case with loss proportional to the transition time.

Markov renewal programming problems are a fertile source of large-scale linear programs. Many problems that look deceptively simple at first sight can lead to linear programs with hundreds or thousands of constraints because of the detailed state description required to make all decision points regeneration points. But often we need not throw in the towel. Another look generally reveals that the constraint matrix is sparse and structured so as to be amenable to decomposition.

XI. ESTIMATION, INFERENCE, AND ADAPTIVE CONTROL

Moore and Pyke [35] develop estimators for the $\{p_{ij}\}$ and the $\{F_{ij}\}$ and their large sample distributions. For statistical inference in birth and death queueing models, see Wolff [54]. Both of the foregoing approaches are objectivist, i.e., non-Bayesian. When a large number of observations are at hand, the objectivist approach is unobjectionable and difficulties stemming from a possible lack of consensus of prior belief do not emerge. On the other hand, when the observations are few or nonexistent, as is common, a Bayesian approach incorporating prior beliefs and loss functions is essential.[12] Such an approach may be formal or may simply consist of a sensitivity analysis with the outcomes being given subjective weights. In the realm of decision making, policies should adapt to modified beliefs as more observations are taken.

Unfortunately, when the tradeoff between information acquisition and immediate losses is explicitly included in the problem formulation, the number of states generally explodes. Generally explicit inclusion is advisable, because if an average cost criterion is interpreted literally, policies that are absurd for any positive discount rate can result. Thus, from a practical viewpoint, for relatively few problems (see, e.g., [21]) can "optimal" adaptive policies be found; for the remainder, it appears that we must be content with heuristic devices. This area remains largely unexplored and is ripe for investigation.

---

[12] This is, of course, a statement of the author's opinion. These matters are highly controversial.

REFERENCES

1. Barlow, R. E., and F. Proschan, Mathematical Theory of Reliability, John Wiley and Sons, New York, 1965.

2. Bellman, R. E., R. E. Kalaba, and J. Lockett, Numerical Inversion of the Laplace Transform, American Elsevier, New York, 1966.

3. Cheong, C. K., "Geometric Convergence of Semi-Markov Transition Probabilities," Z. Wahrscheinlichkeitstheorie verw. Geb., Vol. 7 (1967), pp. 122-130.

4. Çinlar, E., "Decomposition of a Semi-Markov Process under a Markovian Rule," Austral. J. Statist., Vol. 8 (1966), pp. 163-179.

5. -----, "Decomposition of a Semi-Markov Process under a State Dependent Rule," SIAM J. Appl. Math., Vol. 15 (1967), pp. 252-263.

6. -----, "Time Dependence of Queues with Semi-Markovian Services," J. Appl. Prob., Vol. 4 (1967), pp. 356-364.

7. -----, "Queues with Semi-Markovian Arrivals," J. Appl. Prob., Vol. 4 (1967), pp. 365-379.

8. -----, and R. L. Disney, "Stream of Overflows from a Finite Queue," Operations Research, Vol. 15 (1967), pp. 131-134.

9. Cox, D. R., and H. D. Miller, The Theory of Stochastic Processes, John Wiley and Sons, New York, 1965.

10. Denardo, E. V., "Contraction Mappings in the Theory Underlying Dynamic Programming," SIAM Review, Vol. 9 (1967), pp. 165-177.

11. -----, "Separable Markovian Decision Problems," Management Science, Vol. 14 (1968), pp. 451-462.

12. -----, and B. L. Fox, "Multichain Markov Renewal Programs," SIAM J. Appl. Math., Vol. 16 (1968), pp. 468-487.

13. -----, and L. G. Mitten, "Elements of Sequential Decision Processes," J. of Indust. Engin., Vol. 18 (1967), pp. 106-112.

14. Derman, C., "Denumerable State Markovian Decision Processes--Average Cost Criterion," Ann. Math. Stat., Vol. 37 (1966), pp. 145-1553.

15. Fabens, A. J., "The Solution of Queueing and Inventory Models by Semi-Markov Processes," J. Roy. Statist. Soc. Series B, Vol. 23 (1961), pp. 113-127.

16. Feller, W., An Introduction to Probability Theory and Its Applications, Vol. 2, John Wiley and Sons, New York, 1966.

17. -----, "On Semi-Markov Processes," Proc. Nat. Acad. Sci. USA, Vol. 51 (1964), pp. 653-659.

18. Foley, J. D., "A Markovian Model of the University of Michigan Executive System," Comm. ACM, Vol. 10 (1967), pp. 584-588.

19. Fox, B. L., "Markov Renewal Programming by Linear Fractional Programming," SIAM J. Appl. Math., Vol. 14 (1966), pp. 1418-1432.

20. -----, "Existence of Stationary Optimal Policies for Some Markov Renewal Programs," SIAM Review, Vol. 9 (1967), pp. 573-576.

21. -----, "Adaptive Age Replacement," J. Math. Anal. Appl., Vol. 18 (1967), pp. 365-376.

22. -----, (g, w) - Optima in Markov Renewal Programs, P-3624, The RAND Corporation, Santa Monica, California, July 1967. To appear in Management Science.

23. -----, and D. M. Landi, An Algorithm for Identifying the Ergodic Subchains and Transient States of a Stochastic Matrix, RM-5269-PR, The RAND Corporation, Santa Monica, California, March 1967. To appear in Comm. ACM.

24. -----, and L. W. Miller, Equivalent Stationary Measures for the M/G/1 Queue, P-3590, The RAND Corporation, Santa Monica, California, April 1967.

25. Gaver, D. P., "Observing Stochastic Processes and Approximate Transform Inversion," Operations Research, Vol. 14 (1966), pp. 444-459.

26. Janssen, J., "Processus de renouvellements Markoviens et processus semi-markoviens," Cahiers du C.E.R.O.: Part 1, Vol. 6 (1964), pp. 81-105; Part 2, Vol. 7 (1965), pp. 126-141.

27. Jewell, W. S., "The Properties of Recurrent Event Processes," Operations Research, Vol. 8 (1960), pp. 446-472.

28. -----, "Markov-Renewal Programming, I and II," Operations Research, Vol. 11 (1963), pp. 938-971.

29. -----, "A Note on Markov Renewal Programming," ORC 66-41, Operations Research Center, University of California, Berkeley, November 1966.

30. -----, "Fluctuations of a Renewal Reward Process," J. Math. Anal. Appl., Vol. 19 (1967), pp. 309-329.

31. Khintchine, A. J., "Mathematisches über die Erwartung vor einem öffentlichen Schalter," Mat. Sbornik, Vol. 39 (1932), pp. 73-84 (in Russian; German summary).

32. Miller, B. L., Finite State Continuous Time Markov Decision Processes with a Finite Planning Horizon, Vol. 6 (1968), pp. 266-280.

33. -----, Finite State Continuous Time Markov Decision Processes with an Infinite Planning Horizon, RM-5425-PR, The RAND Corporation, September 1967. To appear in J. Math. Anal. Appl.

34. Miller, M. K., and W. T. Guy, "Numerical Inversion of the Laplace Transform by Use of Jacobi Polynomials," SIAM J. Numer. Analysis, Vol. 3 (1966), pp. 624-635.

35. Moore, E. H., and R. Pyke, "Estimation of the Transition Distributions of a Markov Process," D1-82-0371, Boeing Scientific Research Laboratories, Seattle, Washington, September 1964.

36. Neuts, M. F., "Generating Functions for Markov Renewal Processes," Ann. Math. Stat., Vol. 35 (1964), pp. 431-434.

37. -----, "The Busy Period of a Queue with Batch Service," Operations Research, Vol. 13 (1965), pp. 815-819.

38. -----, "The Single Server Queue with Poisson Input and Semi-Markov Service Times," J. Appl. Prob., Vol. 3 (1966), pp. 202-230.

39. -----, "Semi-Markov Analysis of a Bulk Queue," Bull. Soc. Math. Belgique, Vol. 18 (1966), pp. 28-42.

40. -----, "A General Class of Queues with Poisson Input," Ann. Math. Stat., Vol. 38 (1967), pp. 759-770.

41. -----, "Two Queues in Series with a Finite Intermediate Waiting Room," Report 110, Statistics Department, Purdue University, Lafayette, Indiana, 1967.

42. -----, "A Working Bibliography on Markov Renewal Processes and their Applications," Report 140, Statistics Department, Purdue University, Lafayette, Indiana, 1968.

43. Perrin, E. B., and M. C. Sheps, "Human Reproduction: A Stochastic Process," Biometrics, Vol. 20 (1964), pp. 28-45.

44. Pyke, R., "Markov Renewal Processes: Definitions and Preliminary Properties," Ann. Math. Stat., Vol. 32 (1961), pp. 1231-1242.

45. -----, "Markov Renewal Processes with Finitely Many States," Ann. Math. Stat., Vol. 32 (1961), pp. 1243-1259.

46. -----, and R. Schaufele, "Limit Theorems for Markov Renewal Processes," Ann. Math. Stat., Vol. 35 (1964), pp. 1746-1764.

47. -----, and -----, "The Existence and Uniqueness of Stationary Measures for Markov Renewal Processes," Ann. Math. Stat., Vol. 37 (1966), pp. 1439-1462.

48. Ross, S. M., "Non-discounted Denumerable Markovian Decision Models," Ann. Math. Stat., Vol. 39 (1968), pp. 412-423.

49. Spinelli, R. A., "Numerical Inversion of a Laplace Transform," SIAM J. Numer. Analysis, Vol. 3 (1966), pp. 636-649.

50. Stone, L. D., "On the Distribution of the Maximum of a Semi-Markov Process," Ann. Math. Stat., Vol. 39 (1968), pp. 947-956.

51. Teugels, J. L., "On the Rate of Convergence in Renewal and Markov Renewal Processes," Report 107, Statistics Department, Purdue University, Lafayette, Indiana, 1967.

52. Weiss, G., and M. Zelen, A Stochastic Model for Clinical Trials", (abstract), Biometrics, Vol. 18 (1962), p. 627.

53. Widder, D. V., The Laplace Transform, Princeton University Press, Princeton, 1946.

54. Wolff, R. W., "Problems of Statistical Inference for Birth and Death Queueing Models," Operations Research, Vol. 13 (1965), pp. 343-357.

55. Yackel, J., "Limit Theorems for semi-Markov Processes," Trans. Amer. Math. Soc., Vol. 123 (1966), pp. 402-424.